

Dimension Independent Similarity Computation

Reza Bosagh Zadeh

*Institute for Computational and Mathematical Engineering
Stanford University
Stanford, CA 94305, USA*

REZAB@STANFORD.EDU

Ashish Goel

*Department of Management Science and Engineering
Stanford University
Stanford, CA 94305, USA*

ASHISHG@STANFORD.EDU

Editor: Inderjit Dhillon

Abstract

We present a suite of algorithms for Dimension Independent Similarity Computation (DISCO) to compute all pairwise similarities between very high-dimensional sparse vectors. All of our results are provably independent of dimension, meaning that apart from the initial cost of trivially reading in the data, all subsequent operations are independent of the dimension; thus the dimension can be very large. We study Cosine, Dice, Overlap, and the Jaccard similarity measures. For Jaccard similarity we include an improved version of MinHash. Our results are geared toward the MapReduce framework. We empirically validate our theorems with large scale experiments using data from the social networking site Twitter. At time of writing, our algorithms are live in production at twitter.com.

Keywords: cosine, Jaccard, overlap, dice, similarity, MapReduce, dimension independent

1. Introduction

Computing similarity between all pairs of vectors in large-scale data sets is a challenge. Traditional approaches of sampling the data set are limited and linearly dependent on the dimension of the data. We present an approach whose complexity is independent of the data dimension and geared towards modern distributed systems, in particular the MapReduce framework (Dean and Ghemawat, 2008).

MapReduce is a programming model for processing large data sets, typically used to do distributed computing on clusters of commodity computers. With large amount of processing power at hand, it is very tempting to solve problems by brute force. However, we show how to combine clever sampling techniques with the power of MapReduce to extend its utility.

Consider the problem of finding all pairs of similarities between D indicator (0/1 entries) vectors, each of dimension N . In particular we focus on cosine similarities between all pairs of D vectors in \mathbb{R}^N . Further assume that each dimension is L -sparse, meaning each dimension has at most L non-zeros across all points. For example, typical values to compute similarities between all pairs of a subset of Twitter users can be: $N = 10^9$ (the universe of Twitter users), $D = 10^7$ (a subset of Twitter users), $L = 1000$. To be more specific: each of D users can be represented by an N dimensional vector that indicates which of the N universal users are followed. The sparsity parameter ($L = 1000$) here assumes that each user follows at most 1000 other users.

There are two main complexity measures for MapReduce: “shuffle size” and “reduce-key complexity”, defined shortly (Goel and Munagala, 2012). It can be easily shown that the naive approach for computing similarities will have $O(NL^2)$ emissions, which for the example parameters we gave is infeasible. The number of emissions in the map phase is called the “shuffle size”, since that data needs to be shuffled around the network to reach the correct reducer. Furthermore, the maximum number of items reduced to a single key can be as large as N . Thus the “reduce-key complexity” for the naive scheme is N .

We can drastically reduce the shuffle size and reduce-key complexity by some clever sampling with the DISCO scheme described in this paper. In this case, the output of the reducers are random variables whose expectations are the similarities. Two proofs are needed to justify the effectiveness of this scheme: first, that the expectations are indeed correct and obtained with high probability, and second, that the shuffle size is greatly reduced. We prove both of these claims. In particular, in addition to correctness, we prove that to estimate similarities above ϵ , the shuffle size of the DISCO scheme is only $O(DL \log(D)/\epsilon)$, with no dependence on the dimension N , hence the name.

This means as long as there are enough mappers to read the data, one can use the DISCO sampling scheme to make the shuffle size tractable. Furthermore, each reduce key gets at most $O(\log(D)/\epsilon)$ values, thus making the reduce-key complexity tractable too. Within Twitter inc, we use the DISCO sampling scheme to compute similar users. We have also used the scheme to find highly similar pairs of words by taking each dimension to be the indicator vector that signals in which tweets the word appears. We empirically verify the proven claims and observe large reductions in shuffle size in this paper.

Our sampling scheme can be used to implement many other similarity measures. We focus on four similarity measures: Cosine, Dice, Overlap, and the Jaccard similarity measures. For Jaccard similarity we present an improved version of the well known MinHash scheme (Broder, 1997). Our framework operates under the following assumptions, each of which we justify:

First, we focus on the case where each dimension is sparse and therefore the natural way to store the data is by segmenting it into dimensions. In our application each dimension is represented by a tweet, thus this assumption was natural. Second, our sampling scheme requires a “background model”, meaning the magnitude of each vector is assumed to be known and loaded into memory. In our application this was not a hurdle, since the magnitudes of vectors in our corpus needed to be computed for other tasks. In general this may require another pass over the data set. To further address the issue, in the streaming computation model, we can remove the dependence by paying an extra logarithmic factor in memory used. Third, we prove results on highly similar pairs, since common applications require thresholding the similarity score with a high threshold value.

A ubiquitous problem is finding all pairs of objects that are in some sense similar and in particular more similar than a threshold. For such applications of similarity, DISCO is particularly helpful since higher similarity pairs are estimated with provably better accuracy. There are many examples, including

- Advertiser keyword suggestions: When targeting advertisements via keywords, it is useful to expand the manually input set of keywords by other similar keywords, requiring finding all keywords more similar than a high threshold (Regelson and Fain, 2006). The vector representing a keyword will often be an indicator vector indicating in which dimensions the keywords appear.

- Collaborative filtering: Collaborative filtering applications require knowing which users have similar interests. For this reason, given a person or user, it is required to find all other objects more similar than a particular threshold, for which our algorithms are well suited.
- Web Search: Rewriting queries given to a search engine is a common trick used to expand the coverage of the search results (Abhishek and Hosanagar, 2007). Adding clusters of similar queries to a given query is a natural query expansion strategy.

These applications have been around for many years, but the scale at which we need to solve them keeps steadily increasing. This is particularly true with the rise of the web, and social networks such as Twitter, Facebook, and Google Plus among others. Google currently holds an index of more than 1 trillion webpages, on which duplicate detection must be a daunting task. For our experiments we use the Twitter data set, where the number of tweets in a single day surpasses 200 million. Many of the applications, including ours involve domains where the dimensionality of the data far exceeds the number of points. In our case the dimensionality is large (more than 200 million), and we can prove that apart from initially reading the data (which cannot be avoided), subsequent MapReduce operations will have input size independent of the number of dimensions. Therefore, as long as the data can be read in, we can ignore the number of dimensions.

A common technique to deal with large data sets is to simply sample. Indeed, our approach too, is a sampling scheme; however, we sample in a nontrivial way so that points that have nonzero entries in many dimensions will be sampled with lower probability than points that are only present in a few dimensions. Using this idea, we can remove the dependence on dimensionality while being able to mathematically prove—and empirically verify—accuracy.

Although we use the MapReduce (Dean and Ghemawat, 2008) framework and discuss shuffle size, the sampling strategy we use can be generalized to other frameworks. We focus on MapReduce because it is the tool of choice for large scale computations at Twitter, Google, and many other companies dealing with large scale data. However, the DISCO sampling strategy is potentially useful whenever one is computing a number between 0 and 1 by taking the ratio of an unknown number (the dot product in our case) by some known number (e.g., the magnitude). This is a high-level description, and we give four concrete examples, along with proofs, and experiments. Furthermore, DISCO improves the implementation of the well-known MinHash (Broder, 1997) scheme in any MapReduce setup.

2. Formal Preliminaries

Let $T = \{t_1, \dots, t_N\}$ represent N documents, each of a length of no more than L words. In our context, the documents are tweets (documents are also the same as dimensions for us) from the social networking site Twitter, or could be fixed-window contexts from a large corpus of documents, or any other data set where dimensions have no more than L nonzero entries and the data set is available dimension-by-dimension. We are interested in similarity scores between pairs of words in a dictionary containing D words $\{w_1, \dots, w_D\}$. The number of documents in which two words w_i and w_j co-occur is denoted $\#(w_i, w_j)$. The number of documents in which a single word w_i occurs is denoted $\#(w_i)$.

To each word in the dictionary, an N -dimensional indicator vector is assigned, indicating in which documents the word occurs. We operate within a MapReduce framework where each document is an input record. We denote the number of items output by the map phase as the ‘shuffle size’.

Our algorithms will have shuffle sizes that provably do not depend on N , making them particularly attractive for very high dimensional tasks.

We focus on 5 different similarity measures, including cosine similarity, which is very popular and produces high quality results across different domains (Chien and Immorlica, 2005; Chuang and Chien, 2005; Sahami and Heilman, 2006; Spertus et al., 2005). Cosine similarity is simply the vector normalized dot product: $\frac{\#(x,y)}{\sqrt{\#(x)}\sqrt{\#(y)}}$ where $\#(x) = \sum_{i=1}^N x[i]$ and $\#(x,y) = \sum_{i=1}^N x[i]y[i]$. In addition to cosine similarity, we consider many variations of similarity scores that use the dot product. They are outlined in Table 1.

To compare the performance of algorithms in a MapReduce framework, we report and analyze shuffle size, which is more reliable than wall time or any other implementation-specific measure. We define shuffle size as the total output of the Map Phase, which is what will need to be “shuffled” before the Reduce phase can begin. Our results and theorems hold across any MapReduce implementation such as Hadoop (Borthakur, 2007; Gates et al., 2009) or Google’s MapReduce (Dean and Ghemawat, 2008). We focus on shuffle size because after one trivially reads in the data via mappers, the shuffle phase is the bottleneck since our mappers and reducers are all linear in their input size. In general, MapReduce algorithms are usually judged by two performance measures: largest reduce bucket and shuffle size. Since the largest reduce bucket is not at all a problem for us, we focus on shuffle size.

For many applications, and our applications in particular, input vectors are sparse, meaning that the large majority of entries are 0. A sparse vector representation for a vector x is the set of all pairs $(i, x[i])$ such that $x[i] > 0$ over all $i = 1 \dots N$. The size of a vector x , which we denote as $\#(x)$ is the number of such pairs. We focus on the case where each *dimension* is sparse and therefore the natural way to store the data is segmented into dimensions.

Throughout the paper we formally prove results for pairs more similar than a threshold, called ϵ . Our algorithms will often have a tradeoff between accuracy and shuffle size, where the tradeoff parameter is p . By design, the larger p becomes, the more accurate the estimates and the larger the shuffle size becomes.

We assume that the dictionary can fit into memory but that the number of dimensions (documents) is so large that many machines will be needed to even hold the documents on disk. Note that documents and dimensions are the same thing, and we will use these terms interchangeably. We also assume that the magnitudes of the dimensions are known and available in all mappers and reducers.

Similarity	Definition	Shuffle Size	Reduce-key size
Cosine	$\frac{\#(x,y)}{\sqrt{\#(x)}\sqrt{\#(y)}}$	$O(DL\log(D)/\epsilon)$	$O(\log(D)/\epsilon)$
Jaccard	$\frac{\#(x,y)}{\#(x)+\#(y)-\#(x,y)}$	$O((D/\epsilon)\log(D/\epsilon))$	$O(\log(D/\epsilon)/\epsilon)$
Overlap	$\frac{\#(x,y)}{\min(\#(x),\#(y))}$	$O(DL\log(D)/\epsilon)$	$O(\log(D)/\epsilon)$
Dice	$\frac{2\#(x,y)}{\#(x)+\#(y)}$	$O(DL\log(D)/\epsilon)$	$O(\log(D)/\epsilon)$

Table 1: Similarity measures and the bounds we obtain. All sizes are independent of N , the dimension. These are bounds for shuffle size without combining. Combining can only bring down these sizes.

We are interested in several similarity measures outlined in Table 1. For each, we prove the shuffle size for the bottleneck step of the pipeline needed to compute all nonzero scores. Our goal is to compute similarity scores between all pairs of words, and prove accuracy results for pairs which have similarity above ϵ . The naive approach is to first compute the dot product between all pairs of words in a Map-Reduce (Dean and Ghemawat, 2008) style framework. Mappers act on each document, and emit key-value pairs of the form $(w_i, w_j) \rightarrow 1$. These pairs are collected with the sum reducer, which gives the dot product $\#(w_i, w_j)$. The product is then used to trivially compute the formulas in Table 1. The difficulty of computing the similarity scores lies almost entirely in computing $\#(w_i, w_j)$.

Shuffle size is defined as the number of key-value pairs emitted. The shuffle size for the above naive approach is $N\binom{L}{2} = O(NL^2)$, which can be prohibitive for N too large. Our algorithms remove the dependence on N by dampening the number of times popular words are emitted as part of a pair (e.g., see algorithm 5). Instead of emitting every pair, only some pairs are output, and instead of computing intermediary dot products, we directly compute the similarity score.

3. Complexity Measures

An in-depth discussion of complexity measures for MapReduce is available Goel and Munagala (2012). We focus on shuffle size and reduce-key complexity for two reasons. First, focusing on reduce key complexity captures the performance of an idealized Map-Reduce system with infinitely many mappers and reducers and no coordination overheads. A small reduce-key complexity guarantees that a Map-Reduce algorithm will not suffer from “the curse of the last reducer” (Suri and Vassilvitskii, 2011), a phenomenon in which the average work done by all reducers may be small, but due to variation in the size of reduce records, the total wall clock time may be extremely large or, even worse, some reducers may run out of memory.

Second, we focus on shuffle size since the total size of input/output by all mappers/reducers captures the total file I/O done by the algorithm, and is often of the order of the shuffle size.

Lastly, our complexity measures depend only on the algorithm, not on details of the Map-Reduce installation such as the number of machines, the number of Mappers/Reducers etc., which is a desirable property for the analysis of algorithms (Goel and Munagala, 2012).

4. Related Work

Previously the all-pairs similarity search problem has been studied (Broder et al., 1997; Broder, 1997), in the context of identifying near-duplicate web pages. We describe MinHash later in Section 5.2. We improve the vanilla implementation of MinHash on MapReduce to arrive at better shuffle sizes without (effectively) any loss in accuracy. We prove these results in Section 5.2, and verify them experimentally.

There are many papers discussing the all pairs similarity computation problem. In other work (Elsayed et al., 2008; Campagna and Pagh, 2012), the MapReduce framework is targeted, but there is still a dependence on the dimensionality. Other authors (Lin, 2009) have considered computing all pairs on MapReduce, but there is a focus on the life sciences domain. The all-pairs similarity search problem has also been addressed in the database community, where it is known as the similarity join problem (Arasu et al., 2006; Chaudhuri et al., 2006; Sarawagi and Kirpal, 2004). These papers are loosely assigned to one of two categories: First, signatures of the points are obtained to convert the nonexact matching problem to an exact matching problem, followed by hash-map and filtering false positives. Second, inverted list based solutions exploit information retrieval techniques.

In another paper Baraglia et al. (2010), the authors consider the problem of finding highly similar pairs of documents in MapReduce. They discuss shuffle size briefly but do not provide proven guarantees. They also discuss reduce key size, but again do not provide bounds on the size.

There is large body of work on the nearest neighbors problem, which is the problem of finding the k nearest neighbors of a given query point (Charikar, 2002; Fagin et al., 2003; Gionis et al., 1999; Indyk and Motwani, 1998). Our problem, the all-pairs similarities computation, is related, but not identical to the nearest neighbor problem. As one of the reviewers astutely pointed out, when only the k nearest neighbors need to be found, optimizations can be made for that. In particular, there are two problems with pairwise similarity computation, large computation time with respect to dimensions, which we address, and large computation time with respect to the number of points.

Some researchers (Pantel et al., 2009) proposed a highly scalable term similarity algorithm, implemented in the MapReduce framework, and deployed an over 200 billion word crawl of the Web to compute pairwise similarities between terms. Their results are still dependent upon the dimension and they only provide experimental evaluations of their algorithms.

Other related work includes clustering of web data (Beeferman and Berger, 2000; Chien and Immorlica, 2005; Sahami and Heilman, 2006; Spertus et al., 2005). These applications of clustering typically employ relatively straightforward exact algorithms or approximation methods for computing similarity. Our work could be leveraged by these applications for improved performance or higher accuracy through reliance on proven results. We have also used DISCO in finding similar users in a production environment Gupta et al. (2013).

5. Algorithms and Methods

The Naive algorithm for computing similarities first computes dot products, then simply divides by whatever is necessary to obtain a similarity score. I.e. in a MapReduce implementation:

1. Given document t , Map using NaiveMapper (Algorithm 1)
2. Reduce using the NaiveReducer (Algorithm 2)

Algorithm 1 NaiveMapper(t)

for all pairs (w_1, w_2) in t **do**
 emit $((w_1, w_2) \rightarrow 1)$
end for

Algorithm 2 NaiveReducer($(w_1, w_2), \langle r_1, \dots, r_R \rangle$)

$a = \sum_{i=1}^R r_i$
output $\frac{a}{\sqrt{\#(w_1)\#(w_2)}}$

The above steps will compute all dot products, which will then be scaled by appropriate factors for each of the similarity scores. Instead of using the Naive algorithm, we modify the mapper of the naive algorithm and replace it with Algorithm 3, and replace the reducer with Algorithm 4 to directly compute the actual similarity score, rather than the dot products. The following sections detail how to obtain dimensionality independence for each of the similarity scores.

Algorithm 3 DISCOMapper(t)

for all pairs (w_1, w_2) in t **do**
 emit using custom Emit function
end for

Algorithm 4 DISCOReducer($(w_1, w_2), \langle r_1, \dots, r_R \rangle$)

$a = \sum_{i=1}^R r_i$
output $a \frac{\varepsilon}{p}$

5.1 Cosine Similarity

To remove the dependence on N , we replace the emit function with Algorithm 5.

Algorithm 5 CosineSampleEmit(w_1, w_2) - p/ε is the oversampling parameter

With probability

$$\frac{p}{\varepsilon} \frac{1}{\sqrt{\#(w_1)}\sqrt{\#(w_2)}}$$

emit $((w_1, w_2) \rightarrow 1)$

Note that the more popular a word is, the less likely it is to be output. This is the key observation leading to a shuffle size that is independent of the dimension. We use a slightly different reducer, which instead of calculating $\#(w_1, w_2)$, computes $\cos(w_1, w_2)$ directly without any intermediate steps. The exact estimator is given in the proof of Theorem 1.

Since the CosineSampleEmit algorithm is only guaranteed to produce the correct similarity score in expectation, we must show that the expected value is highly likely to be obtained. This guarantee is given in Theorem 1.

Theorem 1 For any two words x and y having $\cos(x, y) \geq \varepsilon$, let $X_1, X_2, \dots, X_{\#(x,y)}$ represent indicators for the coin flip in calls to *CosineSampleEmit* with x, y parameters, and let $X = \sum_{i=1}^{\#(x,y)} X_i$. For any $1 > \delta > 0$, we have

$$\Pr \left[\frac{\varepsilon}{p} X > (1 + \delta) \cos(x, y) \right] \leq \left(\frac{e^\delta}{(1 + \delta)^{(1 + \delta)}} \right)^p$$

and

$$\Pr \left[\frac{\varepsilon}{p} X < (1 - \delta) \cos(x, y) \right] < \exp(-p\delta^2/2)$$

Proof We use $\frac{\varepsilon}{p} X$ as the estimator for $\cos(x, y)$. Note that

$$\mu_{xy} = E[X] = \#(x, y) \frac{p}{\varepsilon} \frac{1}{\sqrt{\#(x)} \sqrt{\#(y)}} = \frac{p}{\varepsilon} \cos(x, y) \geq p$$

Thus by the multiplicative form of the Chernoff bound,

$$\begin{aligned} \Pr \left[\frac{\varepsilon}{p} X > (1 + \delta) \cos(x, y) \right] &= \Pr \left[\frac{\varepsilon}{p} X > (1 + \delta) \frac{\varepsilon}{p} \frac{p}{\varepsilon} \cos(x, y) \right] \\ &= \Pr [X > (1 + \delta) \mu_{xy}] < \left(\frac{e^\delta}{(1 + \delta)^{(1 + \delta)}} \right)^{\mu_{xy}} \leq \left(\frac{e^\delta}{(1 + \delta)^{(1 + \delta)}} \right)^p \end{aligned}$$

Similarly, by the other side of the multiplicative Chernoff bound, we have

$$\Pr \left[\frac{\varepsilon}{p} X < (1 - \delta) \cos(x, y) \right] = \Pr [X < (1 - \delta) \mu_{xy}] < \exp(-\mu_{xy} \delta^2/2) \leq \exp(-p\delta^2/2)$$

■

Since there are $\binom{D}{2}$ pairs of words in the dictionary, set $p = \log(D^2) = 2\log(D)$ and use union bound with theorem 1 to ensure the above bounds hold simultaneously for all pairs x, y having $\cos(x, y) \geq \varepsilon$. It is worth mentioning that $\left(\frac{e^\delta}{(1 + \delta)^{(1 + \delta)}} \right)$ and $\exp(-\delta^2/2)$ are always less than 1; thus raising them to the power of p brings them down exponentially in p . Note that although we decouple p and ε for the analysis of shuffle size, in a good implementation they are tightly coupled, and the expression p/ε can be thought of as a single parameter to trade off accuracy and shuffle size. In a large scale experiment described in section 9.1 we set p/ε in a coupled manner.

Now we show that the shuffle size is independent of N , which is a great improvement over the naive approach when N is large. To see the usefulness of these bounds, it is worth noting that we prove them almost optimal (i.e., no other algorithm can do much better). In Theorem 2 we prove that any algorithm that purports to accurately calculate highly similar pairs must at least output them, and sometimes there are at least DL such pairs, and so any algorithm that is accurate on highly similar pairs must have at least DL shuffle size. We are off optimal here by only a $\log(D)/\varepsilon$ factor.

Theorem 2 The expected shuffle size for *CosineSampleEmit* is $O(DL \log(D)/\varepsilon)$ and $\Omega(DL)$.

Proof The expected contribution from each pair of words will constitute the shuffle size:

$$\begin{aligned}
 & \sum_{i=1}^D \sum_{j=i+1}^D \sum_{k=1}^{\#(w_i, w_j)} \Pr[\text{CosineSampleEmit}(w_i, w_j)] \\
 &= \sum_{i=1}^D \sum_{j=i+1}^D \#(w_i, w_j) \Pr[\text{CosineSampleEmit}(w_i, w_j)] \\
 &= \sum_{i=1}^D \sum_{j=i+1}^D \frac{p}{\varepsilon} \frac{\#(w_i, w_j)}{\sqrt{\#(w_i)} \sqrt{\#(w_j)}} \\
 &\leq \frac{p}{2\varepsilon} \sum_{i=1}^D \sum_{j=i+1}^D \#(w_i, w_j) \left(\frac{1}{\#(w_i)} + \frac{1}{\#(w_j)} \right) \\
 &\leq \frac{p}{\varepsilon} \sum_{i=1}^D \frac{1}{\#(w_i)} \sum_{j=1}^D \#(w_i, w_j) \\
 &\leq \frac{p}{\varepsilon} \sum_{i=1}^D \frac{1}{\#(w_i)} L \#(w_i) = \frac{p}{\varepsilon} LD = O(DL \log(D)/\varepsilon).
 \end{aligned}$$

The first inequality holds because of the Arithmetic-Mean-Geometric-Mean inequality applied to $\{1/\#(w_i), 1/\#(w_j)\}$. The last inequality holds because w_i can co-occur with at most $\#(w_i)L$ other words. It is easy to see via Chernoff bounds that the above shuffle size is obtained with high probability.

To see the lower bound, we construct a data set consisting of D/L distinct documents of length L ; furthermore each document is duplicated L times. To construct this data set, consider grouping the dictionary into D/L groups, each group containing L words. A document is associated with every group, consisting of all the words in the group. This document is then repeated L times. In each group, it is trivial to check that all pairs of words have similarity exactly 1. There are $\binom{L}{2}$ pairs for each group and there are D/L groups, making for a total of $(D/L) \binom{L}{2} = \Omega(DL)$ pairs with similarity 1, and thus also at least ε . Since any algorithm that purports to accurately calculate highly-similar pairs must at least *output* them, and there are $\Omega(DL)$ such pairs, we have the lower bound. ■

It is important to observe what happens if the output “probability” is greater than 1. We certainly Emit, but when the output probability is greater than 1, care must be taken during reducing to scale by the correct factor, since it won’t be correct to divide by p/ε , which is the usual case when the output probability is less than 1. Instead, we must divide by $\sqrt{\#(w_1)} \sqrt{\#(w_2)}$ because for the pairs where the output probability is greater than 1, CosineSampleEmit and Emit are the same. Similar corrections have to be made for the other similarity scores (Dice and Overlap, but not MinHash), so we do not repeat this point. Nonetheless it is an important one which arises during implementation.

Finally it is easy to see that the largest reduce key will have at most $p/\varepsilon = O(\log(D)/\varepsilon)$ values.

Theorem 3 *The expected number of values mapped to a single key by DISCOMapper is p/ε .*

Proof Note that the output of DISCOReducer is a number between 0 and 1. Since this is obtained by normalizing the sum of all values reduced to the key by p/ε , and all summands are at most 1, we trivially get that the number of summands is at most p/ε . ■

5.2 Jaccard Similarity

Traditionally MinHash (Broder, 1997) is used to compute Jaccard similarity scores between all pairs in a dictionary. We improve the MinHash scheme to run much more efficiently with a smaller shuffle size.

Let $h(t)$ be a hash function that maps documents to distinct numbers in $[0, 1]$, and for any word w define $g(w)$ (called the MinHash of w) to be the minimum value of $h(t)$ over all t that contain w . Then $g(w_1) = g(w_2)$ exactly when the minimum hash value of the union with size $\#(w_1) + \#(w_2) - \#(w_1, w_2)$ lies in the intersection with size $\#(w_1, w_2)$. Thus

$$\Pr[g(w_1) = g(w_2)] = \frac{\#(w_1, w_2)}{\#(w_1) + \#(w_2) - \#(w_1, w_2)} = \text{Jac}(w_1, w_2)$$

Therefore the indicator random variable that is 1 when $g(w_1) = g(w_2)$ has expectation equal to the Jaccard similarity between the two words. Unfortunately the variable has too high a variance to be useful on its own. The idea of the MinHash scheme is to reduce the variance by averaging together k of these variables constructed in the same way with k different hash functions. We index these k functions using the notation $g_j(w)$ to denote the MinHash of hash function $h_j(t)$. We denote the computation of hashes as ‘MinHashMap’. Specifically, MinHashMap is defined as Algorithm 6.

To estimate $\text{Jac}(w_1, w_2)$ using this version of the scheme, we simply count the number of hash functions for which $g(w_1) = g(w_2)$, and divide by k to get an estimate of $\text{Jac}(w_1, w_2)$. By the multiplicative Chernoff bound for sums of 0-1 random variables as seen in Theorem 1, setting $k = O(1/\epsilon)$ will ensure that w.h.p. a similarity score that is above ϵ has a relative error of no more than δ . Qualitatively, this theorem is the same as that given in the original MinHash paper (Broder, 1997) (where MinHash is introduced), and we do not claim the following as a new contribution; however, we include it for completeness. More rigorously,

Theorem 4 Fix any two words x and y having $\text{Jac}(x, y) \geq \epsilon$. Let X_1, X_2, \dots, X_k represent indicators for $\{g_1(x) = g_1(y), \dots, g_k(x) = g_k(y)\}$ and $X = \sum_{i=1}^k X_i$. For any $1 > \delta > 0$ and $k = c/\epsilon$, we have

$$\Pr[X/k > (1 + \delta)\text{Jac}(x, y)] \leq \left(\frac{e^\delta}{(1 + \delta)^{(1+\delta)}} \right)^c$$

and

$$\Pr[X/k < (1 - \delta)\text{Jac}(x, y)] \leq \exp(-c\delta^2/2)$$

Proof We use X/k as the estimator for $\text{Jac}(x, y)$. Note that $E[X] = k\text{Jac}(x, y) = (c/\epsilon)\text{Jac}(x, y) \geq c$. Now by standard Chernoff bounds we have

$$\Pr[X/k > (1 + \delta)\text{Jac}(x, y)] = \Pr[X > (1 + \delta)E[X]] \leq \left(\frac{e^\delta}{(1 + \delta)^{(1+\delta)}} \right)^{E[X]} \leq \left(\frac{e^\delta}{(1 + \delta)^{(1+\delta)}} \right)^c$$

Similarly, by the other side of the multiplicative Chernoff bound, we have

$$\Pr[X/k < (1 - \delta)\text{Jac}(x, y)] \leq \exp(-c\delta^2/2)$$

■

The MapReduce implementation of the above scheme takes in documents and for each unique word in the document outputs k hash values.

Algorithm 6 MinHashMap($t, \langle w_1, \dots, w_L \rangle$)

```

for  $i = 1$  to  $L$  do
    for  $j = 1$  to  $k$  do
        emit  $((w_i, j) \rightarrow h_j(t))$ 
    end for
end for
    
```

1. Given document t , Map using MinHashMap (Algorithm 6)
2. Reduce using the min reducer
3. Now that minhash values for all hash functions are available, for each pair we can simply compute the number of hash collisions and divide by the total number of hash functions

Algorithm 7 MinHashSampleMap($t, \langle w_1, \dots, w_L \rangle$)

```

for  $i = 1$  to  $L$  do
    for  $j = 1$  to  $k$  do
        if  $h_j(t) \leq \frac{c \log(Dk)}{\#(w_i)}$  then
            emit  $((w_i, j) \rightarrow h_j(t))$ 
        end if
    end for
end for
    
```

Recall that a document has at most L words. This naive Mapper will have shuffle size $NLk = O(NL/\epsilon)$, which can be improved upon. After the map phase, for each of the k hash functions, the standard MinReducer is used, which will compute $g_j(w)$. These MinHash values are then simply checked for equality. We modify the initial map phase, and prove that the modification brings down shuffle size while maintaining correctness. The modification is seen in algorithm 7, note that c is a small constant we take to be 3.

We now prove that MinHashSampleMap will with high probability Emit the minimum hash value for a given word w and hash function h , thus ensuring the steps following MinHashSampleMap will be unaffected.

Theorem 5 *If the hash functions h_1, \dots, h_k map documents to $[0, 1]$ uniform randomly, and $c = 3$, then with probability at least $1 - \frac{1}{(Dk)^2}$, for all words w and hash functions $h \in \{h_1, \dots, h_k\}$, MinHashSampleMap will emit the document that realizes $g(w)$.*

Proof Fix a word w and hash function h and let $z = \min_{t|w \in t} h(t)$. Now the probability that MinHashSampleMap will not emit the document that realizes $g(w)$ is

$$\Pr \left[z > \frac{c \log(Dk)}{\#(w)} \right] = \left(1 - \frac{c \log(Dk)}{\#(w)} \right)^{\#(w)} \leq e^{-c \log(Dk)} = \frac{1}{(Dk)^c}$$

Thus for a single w and h we have shown `MinHashSampleMap` will w.h.p. emit the hash that realizes the `MinHash`. To show the same result for *all* hash functions and words in the dictionary, set $c = 3$ and use union bound to get a $(\frac{1}{Dk})^2$ bound on the probability of error for *any* w_1, \dots, w_D and h_1, \dots, h_k . ■

Now that we have correctness via theorems 4 and 5, we move on to calculating the shuffle size for `MinHashSampleMap`.

Theorem 6 *MinHashSampleMap* has expected shuffle size $O(Dk \log(Dk)) = O((D/\epsilon) \log(D/\epsilon))$.

Proof Simply adding up the expectations for the emissions indicator variables, we see the shuffle size is bounded by

$$\sum_{h \in \{h_1, \dots, h_k\}} \sum_{w \in \{w_1, \dots, w_D\}} \frac{c \log(Dk)}{\#(w)} \#(w) = Dkc \log(Dk)$$

Setting $c = 3$ and $k = 1/\epsilon$ gives the desired bound. ■

All of the reducers used in our algorithms are associative and commutative operations (sum and min), and therefore can be combined for optimization. Our results do not change qualitatively when combiners are used, except for one case. A subtle point arises in our claim for improving `MinHash`. If we combine with m mappers, then the naive `MinHash` implementation will have a shuffle size of $O(mDk)$ whereas `DISCO` provides a shuffle size of $O(Dk \log(Dk))$. Since m is usually set to be a very large constant (one can easily use 10,000 mappers in a standard Hadoop implementation), removing the dependence on m is beneficial. For the other similarity measures, combining the Naive mappers can only bring down the shuffle size to $O(mD^2)$, which `DISCO` improves upon asymptotically by obtaining a bound of $O(DL/\epsilon \log(D))$ without even combining, so combining will help even more. In practice, `DISCO` can be easily combined, ensuring superior performance both theoretically and empirically.

6. Cosine Similarity in Streaming Model

We briefly depart from the MapReduce framework and instead work in the “Streaming” framework. In this setting, data is streamed dimension-by-dimension through a single machine that can at any time answer queries of the form “what is the similarity between points x and y considering all the input so far?”. The main performance measure is how much memory the machine uses and queries will be answered in constant time. We describe the algorithm only for cosine similarity, but an almost identical algorithm will work for dice, and overlap similarity.

Our algorithm will be very similar to the mapreduce setup, but in place of emitting pairs to be shuffled for a reduce phase, we instead insert them into a hash map H , keyed by pairs of words, with each entry holding a bag of emissions. H is used to track the emissions by storing them in a bag associated with the pair. Since all data streams through a single machine, for any word x , we can keep a counter for $\#(x)$. This will take D counters worth of memory, but as we see in Theorem 1 this memory usage will be dominated by the size of H . Each emission is decorated with the probability of emission. There are two operations to be described: the update that occurs when a new dimension (document) arrives, and the constant time algorithm used to answer similarity queries.

Update. On an update we are given a document. For each pair of words x, y in the document, with independent coin flips of probability $q = \frac{p}{\epsilon} \frac{1}{\sqrt{\#(x)\#(y)}}$ we look up the bag associated with (x, y) in H and insert q into it. It is important to note that the emission is being done with probability q and q is computed using the *current values* of $\#(x)$ and $\#(y)$. Thus if a query comes after this update, we must take into account all new information. This is done via subsampling and will be explained shortly. It remains to show how to use these probabilities to answer queries.

Query. We now describe how to answer the only query. Let the query be for the similarity between words x and y . Recall at this point we know both $\#(x)$ and $\#(y)$ exactly, for the data seen so far. We look up the pair (x, y) in H , and grab the associated bag of emissions. Recall from above that emission is decorated with the probability of emission q_i for the i 'th entry in the bag. Unfortunately q_i will be larger than we need it to be, since it was computed at a previous time, when there had been fewer occurrences of x and y . To remedy this, we independently subsample each of the emissions for the pair x, y with coin flips of probability $\frac{p}{\epsilon} \frac{1}{q_i \sqrt{\#(x)\#(y)}}$. For each pair x, y seen in the input, there will be exactly

$$q_i \frac{p}{\epsilon} \frac{1}{q_i \sqrt{\#(x)\#(y)}} = \frac{p}{\epsilon} \frac{1}{\sqrt{\#(x)\#(y)}}$$

probability of surviving emission *and* the subsampling. Finally, since the pair x, y is seen exactly $\#(x, y)$ times, the same estimator used in Theorem 2 will have expectation equal to $\cos(x, y)$. Furthermore, since before subsampling we output in expectation more pairs than `CosineSampleEmit`, the Chernoff bound of Theorem 2 still holds. Finally, to show that H cannot grow too large, we bound its size in Theorem 7.

Theorem 7 *The streaming algorithm uses at most $O(DL \lg(N) \log(D)/\epsilon)$ memory.*

Proof We only need to bound the size of H . Consider a word x and all of its occurrences in documents $t_1, \dots, t_{\#(x)}$ at final time (i.e., after all N documents have been processed). We conceptually and only for this analysis construct a new larger data set C' where each word x is removed and its occurrences are replaced *in order* with $\lceil \lg \#(x) \rceil + 1$ new words $x_1, \dots, x_{\lceil \lg \#(x) \rceil + 1}$, so we are effectively segmenting (in time) the occurrences of x . With this in mind, we construct C' so that each x_i will replace 2^{i-1} occurrences of x , in time order. That is, we will have $\#(x_i) = 2^{i-1}$.

Note that our streaming algorithm updates the counters for words with every update. Consider what happens if instead of updating every time, the streaming algorithm somehow in advance knew and used the final $\#(x)$ values after all documents have been processed. We call this the “all-knowing” version. The size of H for such an all-knowing algorithm is the same as the shuffle size for the DISCO sampling scheme analyzed in Theorem 2, simply because there is a bijection between the emits of `CosineSampleEmit` and inserts into H with exactly the same coin flip probability. We now use this observation and C' .

We show that the memory used by H when our algorithm is run on C' with the all-knowing counters dominates (in expectation) the size of H for the original data set in the streaming model, thus achieving the claimed bound in the current theorem statement.

Let $\text{PrHashMapInsert}(x, y)$ denote the probability of inserting the pair x, y when we run the all-knowing version of the streaming algorithm on C' . Let $\text{PrStreamEmit}(x, y, a, b)$ denote the probability of emitting the pair x, y in the streaming model with input C , after observing x, y exactly a, b

times, respectively. With these definitions we have

$$\begin{aligned}
 \text{PrStreamEmit}(x, y, a, b) &= \frac{p}{\epsilon} \frac{1}{\sqrt{ab}} \\
 &\leq \frac{p}{\epsilon} \frac{1}{2^{\lfloor \lg a \rfloor} 2^{\lfloor \lg b \rfloor}} \leq \frac{p}{\epsilon} \frac{1}{\sqrt{\#(x_{\lfloor \lg a \rfloor})} \sqrt{\#(x_{\lfloor \lg b \rfloor})}} \\
 &= \text{PrHashMapInsert}(x_{\lfloor \lg a \rfloor + 1}, y_{\lfloor \lg b \rfloor + 1}).
 \end{aligned}$$

The first inequality holds by properties of the floor function. The second inequality holds by definition of C' . The dictionary size for C' is $O(D \lg(N))$ where D is the original dictionary size for C . Using the same analysis used in Theorem 2, the shuffle size for C' is at most $O(DL \lg(N) \log(D)/\epsilon)$ and therefore so is the size of H for the all-knowing algorithm run on C' , and, by the analysis above, so is the hash map size for the original data set C . \blacksquare

7. Correctness and Shuffle Size Proofs for other Similarity Measures

We now give consideration to similarity scores other than cosine similarity, along proofs of their shuffle size.

7.1 Overlap Similarity

Overlap similarity follows the same pattern we used for cosine similarity; thus we only explain the parts that are different. The emit function changes to Algorithm 8.

Algorithm 8 $\text{OverlapSampleEmit}(w_1, w_2) - p/\epsilon$ is the oversampling parameter

With probability

$$\frac{p}{\epsilon} \frac{1}{\min(\#(w_1), \#(w_2))}$$

emit $((w_1, w_2) \rightarrow 1)$

The correctness proof is nearly identical to that of cosine similarity so we do not restate it. The shuffle size for OverlapSampleEmit is given by the following theorem.

Theorem 8 *The expected shuffle size for OverlapSampleEmit is $O(DL \log(D)/\epsilon)$.*

Proof The expected contribution from each pair of words will constitute the shuffle size:

$$\begin{aligned}
 &\sum_{i=1}^D \sum_{j=i+1}^D \frac{p}{\epsilon} \frac{\#(w_i, w_j)}{\min(\#(w_i), \#(w_j))} \\
 &\leq \frac{p}{\epsilon} \sum_{i=1}^D \sum_{j=i+1}^D \#(w_i, w_j) \left(\frac{1}{\#(w_i)} + \frac{1}{\#(w_j)} \right) \\
 &\leq \frac{2p}{\epsilon} \sum_{i=1}^D \frac{1}{\#(w_i)} \sum_{j=1}^D \#(w_i, w_j)
 \end{aligned}$$

$$\leq \frac{2p}{\epsilon} \sum_{i=1}^D \frac{1}{\#(w_i)} L\#(w_i) = \frac{2p}{\epsilon} LD = O(DL \log(D)/\epsilon)$$

The first inequality holds trivially. The last inequality holds because w_i can co-occur with at most $\#(w_i)L$ other words. It is easy to see via Chernoff bounds that the above shuffle size is obtained with high probability. ■

7.2 Dice Similarity

Dice similarity follows the same pattern as we used for cosine similarity, thus we only explain the parts that are different. The emit function changes to Algorithm 9.

Algorithm 9 DiceSampleEmit(w_1, w_2) - p/ϵ is the oversampling parameter

With probability

$$\frac{p}{\epsilon} \frac{2}{\#(w_1) + \#(w_2)}$$

emit ($(w_1, w_2) \rightarrow 1$)

The correctness proof is nearly identical to cosine similarity so we do not restate it. The shuffle size for DiceSampleEmit is given by the following theorem:

Theorem 9 *The expected shuffle size for DiceSampleEmit is $O(DL \log(D)/\epsilon)$.*

Proof The expected contribution from each pair of words will constitute the shuffle size:

$$\begin{aligned} & 2 \sum_{i=1}^D \sum_{j=i+1}^D \frac{p}{\epsilon} \frac{\#(w_i, w_j)}{\#(w_i) + \#(w_j)} \\ & \leq \frac{2p}{\epsilon} \sum_{i=1}^D \sum_{j=1}^D \frac{\#(w_i, w_j)}{\#(w_i)} \\ & \leq \frac{2p}{\epsilon} \sum_{i=1}^D \frac{1}{\#(w_i)} L\#(w_i) = \frac{2p}{\epsilon} LD = O(DL \log(D)/\epsilon) \end{aligned}$$

The first inequality holds trivially. The last inequality holds because w_i can co-occur with at most $\#(w_i)L$ other words. It is easy to see via Chernoff bounds that the above shuffle size is obtained with high probability. ■

8. Cross Product

Consider the case in which we have N vectors (and *not* D) in \mathbb{R}^N and wish to compute the cross product of all similarities between a subset D of the N vectors with the remaining $N - D$ vectors. This is different from computing all $\binom{D}{2}$ similarities, which has been the focus of the previous sections.

It is worth a brief mention that the same bounds and almost the same proof of Theorem 2 hold for this case, with one caveat, namely that the magnitudes of all N points must be available to the mappers and reducers. We can mitigate this problem by annotating each dimension with the magnitudes of points beforehand in a separate MapReduce job. However, this MapReduce job will have a shuffle size dependent on N , so we do not consider it a part of the DISCO scheme.

9. Experiments

We use data from the social networking site Twitter, a very popular social network platform. Users interact with Twitter through a web interface, instant messaging clients, or sending mobile text messages. Public updates by users are viewable to the world, and a large majority of Twitter accounts are public. These public tweets provide a large real-time corpus of what is happening in the world.

We have two sets of experiments. The first is a large scale experiment that is in production, to compute similar users in the Twitter follow graph. The second is an experiment to find similar words. The point of the first experiment is show the scalability of our algorithms and the point of the second set of experiments is to show accuracy.

9.1 Similar Users

Consider the problem of finding all pairs of similarities between a subset of $D = 10^7$ twitter users. We would like to find all $\binom{D}{2}$ similarities between these users. The number of dimensions of each user is larger than D , and is denoted $N = 5 \times 10^8$. We used $p/\epsilon = 100$. We have a classifier trained to pick the edges with most interaction and thus can limit the sparsity of the data set with $L = 1000$.

We compute these user similarities daily in a production environment using Hadoop and Pig implementations at Twitter (Olston et al., 2008), with the results supplementing Twitter.com. Using 100 Hadoop nodes we can compute the above similarities in around 2 hours. In this case, the naive scheme would have shuffle size $NL^2 = 5 \times 10^{14}$, which is infeasible but becomes possible with DISCO.

9.2 Similar Words

In the second experiment we used all tweets seen by Twitter.com in one day. The point in the experiment is to find all pairs of similar words, where words are similar if they appear in the same tweets. The number of dimensions $N = 198,134,530$ in our data is equal to the number of tweets and each tweet is a document with size at most 140 characters, providing a small upper bound for L . These documents are ideal for our framework, since our shuffle size upper bound depends on L , which in this case is very small. We used a dictionary of 1000 words advertisers on Twitter are currently targeting to show Promoted Trends, Trends, and Accounts. We also tried a uniformly random sampled dictionary without a qualitative change in results.

The reason we only used $D = 1000$ is because for the purpose of validating our work (i.e., reporting the small errors occurred by our algorithms), we have to compute the true cosine similarities, which means computing true co-occurrence for every pair, which is a challenging task computationally. This was a bottleneck only in *our experiments* for this paper, and does not affect users of our algorithms. We ran experiments with $D = 10^6$, but cannot report true error since finding the true cosine similarities is too computationally intensive. In this regime however, our theorems guarantee that the results are good with high probability.

9.3 Shuffle Size vs Accuracy

We have two parameters to tweak to tradeoff accuracy and shuffle size: ϵ and p . However, since they only occur in our algorithms as the ratio p/ϵ , we simply use that as the tradeoff parameter. The reason we separated p and ϵ was for the theorems to go through nicely, but in reality we only see a single tweaking parameter.

We increase p/ϵ exponentially on the x axis and record the ratio of DISCO shuffle size to the naive implementation. In all cases we can achieve a 90% reduction in shuffle size compared to the naive implementation without sacrificing much accuracy, as see in Figures 2, 4, and 6. The accuracy we report is with respect to true cosine, dice, and overlap similarity.




Figure 1: Average error for all pairs with similarity $\geq \epsilon$. DISCO estimated Cosine error decreases for more similar pairs. Shuffle reduction from naive implementation: 99.39%.

9.4 Error vs Similarity Magnitude

All of our theorems report better accuracy for pairs that have higher similarity than otherwise. To see this empirically, we plot the average error of all pairs that have true similarity above ϵ . These can be seen in Figures 1, 3, 5, 7, and 8. Note that the reason for large portions of the error being constant in these plots is that there are very few pairs with very high similarities, and therefore the error remains constant while ϵ is between the difference of two such very high similarity pairs. Further note that Figure 7 and 8 are so similar because our proposed DISCO MinHash very closely mimics the original MinHash. This is reflected in the strength of the bound in Theorem 5.

10. Conclusions and Future Directions

We presented the DISCO suite of algorithms to compute all pairwise similarities between very high dimensional sparse vectors. All of our results are provably independent of dimension, meaning that apart from the initial cost of trivially reading in the data, all subsequent operations are independent of the dimension and thus the dimension can be very large.

Although we use the MapReduce (Dean and Ghemawat, 2008) and Streaming computation models to discuss shuffle size and memory, the sampling strategy we use can be generalized to




Figure 2: As p/ϵ increases, shuffle size increases and error decreases. There is no thresholding for highly similar pairs here. Ground truth is computed with the naive algorithm. When the true similarity is zero, DISCO always also returns zero, so we always get those pairs right. It remains to estimate those pairs with similarity > 0 , and that is the average relative error for those pairs that we report here.




Figure 3: Average error for all pairs with similarity $\geq \epsilon$. DISCO estimated Dice error decreases for more similar pairs. Shuffle reduction from naive implementation: 99.76%.

other frameworks. We anticipate the DISCO sampling strategy to be useful whenever one is computing a number between 0 and 1 by taking the ratio of an unknown number (the dot product in our case) to some known number (e.g., $\sqrt{\#(x)\#(y)}$ for cosine similarity) (Zadeh and Carlsson, 2013). This is a high-level description, and in order to make it more practical, we give five concrete examples, along with proofs and experiments.




Figure 4: As p/ϵ increases, shuffle size increases and error decreases. There is no thresholding for highly similar pairs here.




Figure 5: Average error for all pairs with similarity $\geq \epsilon$. DISCO estimated Overlap error decreases for more similar pairs. Shuffle reduction from naive implementation: 97.86%.

Acknowledgments

We thank the reviewers for their immensely helpful and thorough comments. We also thank the Twitter Personalization and Recommender Systems team for allowing use of production data. Supported in part by the DARPA xdata program, by grant FA9550-12-1-0411 from the U.S. Air Force Office of Scientific Research (AFOSR) and the Defense Advanced Research Projects Agency (DARPA), and by NSF Award 0915040.




Figure 6: As p/ϵ increases, shuffle size increases and error decreases. There is no thresholding for highly similar pairs here.




Figure 7: Average error for all pairs with similarity $\geq \epsilon$. MinHash Jaccard similarity error decreases for more similar pairs. We are computing error with MinHash here, not ground truth.

References

V. Abhishek and K. Hosanagar. Keyword generation for search engine advertising using semantic similarity between terms. In *EC 2007*, pages 89–94. ACM, 2007.

A. Arasu, V. Ganti, and R. Kaushik. Efficient exact set-similarity joins. In *VLDB 2006*, pages 918–929. VLDB Endowment, 2006.

R. Baraglia, G. De Francisci Morales, and C. Lucchese. Document similarity self-join with mapreduce. In *Data Mining (ICDM), 2010 IEEE 10th International Conference on*, pages 731–736. IEEE, 2010.




Figure 8: Average error for all pairs with similarity $\geq \epsilon$. DISCO MinHash Jaccard similarity error decreases for more similar pairs. We are computing error with MinHash here, not ground truth.

D. Beeferman and A. Berger. Agglomerative clustering of a search engine query log. In *SIGKDD 2000*, pages 407–416. ACM, 2000.

D. Borthakur. The hadoop distributed file system: Architecture and design. *Hadoop Project Website*, 2007.

A.Z. Broder. On the resemblance and containment of documents. In *CCS 1997*, pages 21–29. IEEE, 1997.

A.Z. Broder, S.C. Glassman, M.S. Manasse, and G. Zweig. Syntactic clustering of the web. *Computer Networks and ISDN Systems*, 29(8-13):1157–1166, 1997.

A. Campagna and R. Pagh. Finding associations and computing similarity via biased pair sampling. *Knowledge and Information Systems*, 31(3):505–526, 2012.

M.S. Charikar. Similarity estimation techniques from rounding algorithms. In *STOC 2002*, pages 380–388. ACM, 2002.

S. Chaudhuri, V. Ganti, and R. Kaushik. A primitive operator for similarity joins in data cleaning. In *ICDE 2006*, pages 5–5. IEEE, 2006.

S. Chien and N. Immerlica. Semantic similarity between search engine queries using temporal correlation. In *WWW 2005*, pages 2–11. ACM, 2005.

S.L. Chuang and L.F. Chien. Taxonomy generation for text segments: A practical web-based approach. *ACM TOIS*, 23(4):363–396, 2005.

J. Dean and S. Ghemawat. MapReduce: Simplified data processing on large clusters. *Communications of the ACM*, 51(1):107–113, 2008. ISSN 0001-0782.

- T. Elsayed, J. Lin, and D.W. Oard. Pairwise document similarity in large collections with mapreduce. In *ACL 2008*, pages 265–268. Association for Computational Linguistics, 2008.
- R. Fagin, R. Kumar, and D. Sivakumar. Efficient similarity search and classification via rank aggregation. In *ACM SIGMOD 2003*, pages 301–312. ACM, 2003.
- A.F. Gates, O. Natkovich, S. Chopra, P. Kamath, S.M. Narayanamurthy, C. Olston, B. Reed, S. Srinivasan, and U. Srivastava. Building a high-level dataflow system on top of map-reduce: the pig experience. *Proceedings of the VLDB Endowment*, 2(2):1414–1425, 2009.
- A. Gionis, P. Indyk, and R. Motwani. Similarity search in high dimensions via hashing. In *VLDB 1999*, pages 518–529. Morgan Kaufmann Publishers Inc., 1999.
- A Goel and K Munagala. Complexity measures for map-reduce, and comparison to parallel computing. *Manuscript* <http://www.stanford.edu/~ashishg/papers/mapreducecomplexity.pdf>, 2012.
- P. Gupta, A. Goel, J. Lin, A. Sharma, D. Wang, and R. Zadeh. Wtf: The who to follow service at twitter. *The WWW 2013 Conference*, 2013.
- P. Indyk and R. Motwani. Approximate nearest neighbors: towards removing the curse of dimensionality. In *STOC 1998*, pages 604–613. ACM, 1998.
- J. Lin. Brute force and indexed approaches to pairwise document similarity comparisons with mapreduce. In *SIGIR 2009*, pages 155–162. ACM, 2009.
- C. Olston, B. Reed, U. Srivastava, R. Kumar, and A. Tomkins. Pig latin: a not-so-foreign language for data processing. In *Proceedings of the 2008 ACM SIGMOD International Conference on Management of Data*, pages 1099–1110. ACM, 2008.
- P. Pantel, E. Crestan, A. Borkovsky, A.M. Popescu, and V. Vyas. Web-scale distributional similarity and entity set expansion. In *EMNLP 2009*. ACL, 2009.
- M. Regelson and D. Fain. Predicting click-through rate using keyword clusters. In *Proceedings of the Second Workshop on Sponsored Search Auctions*, volume 9623. Citeseer, 2006.
- M. Sahami and T.D. Heilman. A web-based kernel function for measuring the similarity of short text snippets. In *WWW 2006*, pages 377–386. ACM, 2006.
- S. Sarawagi and A. Kirpal. Efficient set joins on similarity predicates. In *ACM SIGMOD 2004*, pages 743–754. ACM, 2004.
- E. Spertus, M. Sahami, and O. Buyukkokten. Evaluating similarity measures: a large-scale study in the orkut social network. In *SIGKDD 2005*, pages 678–684. ACM, 2005.
- S. Suri and S. Vassilvitskii. Counting triangles and the curse of the last reducer. In *Proceedings of the 20th International Conference on World Wide Web*, pages 607–614. ACM, 2011.
- R.B. Zadeh and G. Carlsson. Dimension independent matrix square using mapreduce. *Symposium on Theory of Computing*, 2013.